

Barrier Bayesian Linear Regression: Online Learning of Control Barrier Conditions for Safety-Critical Control of Uncertain Systems

Lukas Brunke^{1,3,4*}

LUKAS.BRUNKE@ROBOTICS.UTIAS.UTORONTO.CA

Siqi Zhou^{1,3,4*}

SIQI.ZHOU@ROBOTICS.UTIAS.UTORONTO.CA

Angela P. Schoellig^{1,2,3,4}

ANGELA.SCHOELLIG@UTORONTO.CA

¹University of Toronto Institute for Aerospace Studies, Toronto, Ontario, Canada

²Technical University of Munich, Munich, Germany

³University of Toronto Robotics Institute, Toronto, Ontario, Canada

⁴Vector Institute for Artificial Intelligence, Toronto, Ontario, Canada

*Equal contribution

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

In this work, we consider the problem of designing a safety filter for a nonlinear uncertain control system. Our goal is to augment an arbitrary controller with a safety filter such that the overall closed-loop system is guaranteed to stay within a given state constraint set, referred to as being safe. For systems with known dynamics, control barrier functions (CBFs) provide a scalar condition for determining if a system is safe. For uncertain systems, robust or adaptive CBF certification approaches have been proposed. However, these approaches can be conservative or require the system to have a particular parametric structure. For more generic uncertain systems, machine learning approaches have been used to approximate the CBF condition. These works typically assume that the learning module is sufficiently trained prior to deployment. Safety during learning is not guaranteed. We propose a barrier Bayesian linear regression (BBLR) approach that guarantees safe online learning of the CBF condition for the true, uncertain system. We assume that the error between the nominal system and the true system is bounded and exploit the structure of the CBF condition. We show that our approach can safely expand the set of certifiable control inputs despite system and learning uncertainties. The effectiveness of our approach is demonstrated in simulation using a two-dimensional pendulum stabilization task.

Keywords: Safety-critical control, learning control barrier conditions, Bayesian linear regression

1. Introduction

Robots are envisioned to perform increasingly complex tasks in applications ranging from autonomous driving to space exploration. In these applications, robots are required to cope with uncertainties in the operating environment, and their safe operation is critical. In the literature, nonlinear control techniques, such as Lyapunov stability (Khalil, 2002), Hamilton-Jacobi reachability analysis (Mitchell et al., 2005), or control barrier functions (CBFs) (Ames et al., 2019), provide the mathematical foundation for deriving safety conditions when the system dynamics are known. However, it remains challenging to provide safety guarantees if the robot dynamics are uncertain (Brunke

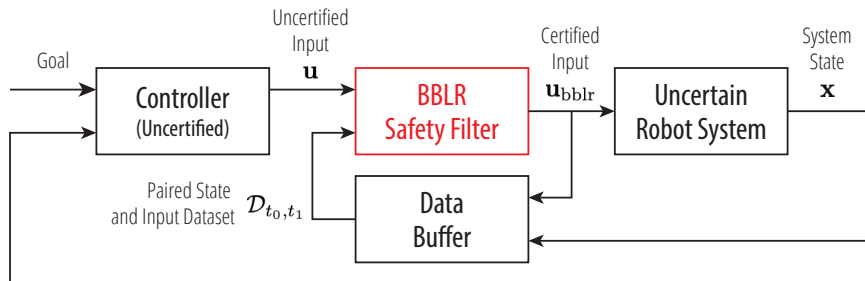


Figure 1: A block diagram of the proposed barrier Bayesian linear regression (BBLR) safety filter approach, which aims to safely learn an uncertain CBF condition online. Augmenting a given, and possibly unsafe controller, with the BBLR safety filter (in red) guarantees that the state of the uncertain robot system stays within a state constraint set.

et al., 2022). In this work, we consider the problem of keeping the system inside of a safe state constraint set. We leverage the CBF certification framework and Bayesian linear regression (BLR) to design an adaptive safety filter that renders any given controller safe despite uncertainties present in the system (Figure 1). The proposed approach exploits the structure of the CBF condition and ensures that an online learned CBF condition is valid by construction.

Inspired by Nagumo’s theorem (Nagumo, 1942), the CBF certification framework provides necessary and sufficient safety conditions for a nonlinear control system to stay within a predefined state constraint set (Ames et al., 2019). In (Ames et al., 2014), CBFs have been used as safety filters to modify unsafe control inputs computed by an arbitrary controller. However, this framework relies on knowing the dynamics of the system. Any uncertainty in the dynamics results in uncertainty in the CBF certification condition, which can lead to violations of the safety constraints. One possibility to account for these uncertainties is by employing robust (Choi et al., 2021) or adaptive (Taylor and Ames, 2020; Lopez et al., 2021) CBF conditions to recover safety. While these approaches guarantee safety, robust approaches are typically overly conservative and adaptive approaches often require an *a priori* known structure of the uncertain dynamics, which can be restrictive in practice.

Recently, learning-based approaches have been developed to address these drawbacks and learn the CBF from data. Compared to approaches that try to learn the system dynamics (Brunke et al., 2022), the advantage of CBF approaches is that they directly learn the error in the CBF condition, which is typically of lower dimension than the full system dynamics (Taylor et al., 2019). In (Choi et al., 2020; Taylor et al., 2020, 2021), neural network models are used to learn the CBF conditions offline from pre-collected data. By representing the CBF condition error by a generic neural network, a larger class of uncertain systems can be considered compared to other adaptive approaches. In (Wang et al., 2018; Ohnishi et al., 2019; Cheng et al., 2019; Fan et al., 2020; Khojasteh et al., 2020), probabilistic learning techniques for CBFs have been proposed to explicitly account for the uncertainty of the learned CBF condition and guarantee safety of the system probabilistically.

However, current learning-based approaches do not guarantee safety while learning the CBF condition, or only do so by assumption (Taylor et al., 2021). This is due to the fact that the learned approximations typically do not have a lower bound that can be conveniently incorporated into the CBF safety analysis. Furthermore, the data collected for training the CBF condition may be subject to errors (e.g., due to numerical differentiation or noise in the system) (Khojasteh et al., 2020) or

may be insufficient for learning the CBF condition well (Taylor et al., 2020, 2021). These errors in the learned CBF condition are not explicitly accounted for in current approaches and can result in unsafe behavior.

In this work, we present an online learning framework, called barrier Bayesian linear regression (BBLR), that exploits the structure of the CBF condition and provably guarantees the satisfaction of the CBF condition *during learning*. Under the assumption that the error between a given nominal system and the true system is bounded, we prove that our proposed BBLR algorithm can expand the set of safe control inputs while keeping the system inside of the safe set. Additionally, our Bayesian approach captures the uncertainty of the learned CBF condition, which allows us to incorporate the uncertainty of the learned CBF condition in the overall safety filter framework. We illustrate the improved performance and safety of our approach on a simulated two-dimensional pendulum.

2. Problem Formulation

We consider the control architecture in Figure 1. The uncertain robot system is modeled as a continuous-time control-affine system with perfect state measurements:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}, \quad (1)$$

where $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^n$ is the state of the system, $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^m$ is the control input and \mathcal{U} is compact, and $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\mathbf{g} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are locally Lipschitz continuous functions. We refer to this system as the true system which is partially unknown.

We consider a nominal control-affine system given by

$$\dot{\hat{\mathbf{x}}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u}, \quad (2)$$

where $\hat{\mathbf{f}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\hat{\mathbf{g}} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are known, locally Lipschitz continuous functions. Without loss of generality, we can write the true dynamics as

$$\dot{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u} + \mathbf{d}(\mathbf{x}, \mathbf{u}), \quad (3)$$

where $\mathbf{d}(\mathbf{x}, \mathbf{u}) = \mathbf{a}(\mathbf{x}) + \mathbf{b}(\mathbf{x})^\top \mathbf{u}$ with $\mathbf{a}(\mathbf{x}) = \mathbf{f}(\mathbf{x}) - \hat{\mathbf{f}}(\mathbf{x})$ and $\mathbf{b}(\mathbf{x}) = \mathbf{g}(\mathbf{x}) - \hat{\mathbf{g}}(\mathbf{x})$. We assume that the difference between the true and nominal system $\mathbf{d}(\mathbf{x}, \mathbf{u})$ is bounded for all $(\mathbf{x}, \mathbf{u}) \in \mathcal{X} \times \mathcal{U}$.

Our goal is to augment a given, not certified state-feedback controller with a safety filter such that the true system's state \mathbf{x} in Equation 1 stays inside a safe set \mathcal{C} , see Figure 1. The safe set \mathcal{C} is assumed to be compact and is defined as the 0-superlevel set of a C^1 function $h : \mathbb{R}^n \rightarrow \mathbb{R}$:

$$\mathcal{C} = \{\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^n \mid h(\mathbf{x}) \geq 0\}. \quad (4)$$

Its boundary is defined as $\partial\mathcal{C} = \{\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^n \mid h(\mathbf{x}) = 0\}$, where for all $\mathbf{x} \in \partial\mathcal{C}$ it holds that $\partial h(\mathbf{x})/\partial \mathbf{x} \neq 0$. We assume that there exists a control input signal $\nu : \mathbb{R}_{\geq 0} \mapsto \mathcal{U}$ such that the true system in Equation 1 stays inside \mathcal{C} . We also assume that there exists a control input signal ν , such that the nominal system in Equation 2 stays inside a subset $\mathcal{C}_{\bar{h}}$ of the safe set \mathcal{C} .

3. Technical Background

In this section, we introduce the background for safely learning control barrier function conditions. We begin this section with the definitions of the extended class \mathcal{K}_∞ comparison functions and positively control invariant sets. We then give an overview of control barrier functions and Bayesian linear regression.

3.1. Definitions

Comparison functions are used to analyze a nonlinear system's stability with Lyapunov functions. CBF conditions rely on comparison functions from the extended class \mathcal{K}_∞ :

Definition 1 (Extended class \mathcal{K}_∞ function) *A continuous function $\gamma : (-b, a) \rightarrow \mathbb{R}$, with $\gamma(0) = 0$, γ is strictly monotonically increasing, and $a, b = \infty$, $\lim_{r \rightarrow \infty} \gamma(r) = \infty$ and $\lim_{r \rightarrow -\infty} \gamma(r) = -\infty$, then γ is said to belong to extended class \mathcal{K}_∞ or equivalently expressed as $\gamma \in \mathcal{K}_{\infty, \epsilon}$.*

In this work, safety and positively control invariance are used synonymously, although, other definitions exist (Brunke et al., 2022). We recall the definition of positively control invariance:

Definition 2 (Positively control invariant set) *Let \mathcal{U} be the set of all $\nu : \mathbb{R}_{\geq 0} \rightarrow \mathcal{U}$. A set $\mathcal{C} \subset \mathcal{X}$ is a positively control invariant set for a control system $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u})$ with $\mathbf{f} : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$ if $\forall \mathbf{x}_0 \in \mathcal{C}$, $\exists \nu \in \mathcal{U}$, $\forall t \in T_{\mathbf{x}_0}^+$, $\phi(t, \mathbf{x}_0, \nu) \in \mathcal{C}$, where $\phi(t, \mathbf{x}_0, \nu)$ is the phase flow of $\mathbf{f}(\mathbf{x}, \mathbf{u})$ starting at \mathbf{x}_0 and following ν , and $T_{\mathbf{x}_0}^+$ is the maximum time interval.*

3.2. Control Barrier Functions

We aim to find a controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ for a control-affine system as in Equation 1 that renders the set \mathcal{C} positively control invariant. Suppose that $\mathbf{u} = \mathbf{k}(\mathbf{x})$ is a state-feedback controller rendering the closed-loop system $\dot{\mathbf{x}} = \mathbf{f}_{\text{cl}}(\mathbf{x}) = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x})$ locally Lipschitz continuous on the domain \mathcal{X} . Then, for any initial condition $\mathbf{x}_0 \in \mathcal{X}$ there exists a maximum interval of existence $T_{\mathbf{x}_0}^+$ such that $\mathbf{x}(t)$ is the unique solution to the closed-loop system on $T_{\mathbf{x}_0}^+$. The definition of a control barrier function following (Ames et al., 2019) is given next.

Definition 3 (Control barrier function (CBF)) *Let $\mathcal{C} \subseteq \mathcal{X} \subset \mathbb{R}^n$ be the superlevel set of a C^1 function $h : \mathcal{D} \rightarrow \mathbb{R}$, then h is a control barrier function (CBF) if there exists an extended class \mathcal{K}_∞ function γ such that for all $\mathbf{x} \in \mathcal{X}$ the control system in Equation 1 satisfies*

$$\sup_{\mathbf{u} \in \mathcal{U}} [L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u}] \geq -\gamma(h(\mathbf{x})). \quad (5)$$

Let $\mathcal{U}_{\text{CBF}}(\mathbf{x}) = \{\mathbf{u} \in \mathcal{U} \mid L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} + \gamma(h(\mathbf{x})) \geq 0\}$ be the set of inputs satisfying the CBF condition. Then, for any $\mathbf{x} \in \mathcal{X}$ any Lipschitz continuous controller $\mathbf{k}(\mathbf{x}) \in \mathcal{U}_{\text{CBF}}(\mathbf{x})$ guarantees positive invariance of \mathcal{C} under the dynamics in Equation 1. Determining the positive invariance of a system amounts to checking if the Lie derivative condition in Equation 5 is satisfied. This is stated in the following theorem (Ames et al., 2017):

Theorem 4 (CBF positive invariant set certification) *Let $\mathcal{C} \subset \mathbb{R}^n$ be defined as in Equation 4 in terms of a C^1 function $h : \mathcal{X} \subset \mathbb{R}^n \rightarrow \mathbb{R}$. If h is a control barrier function on \mathcal{X} and for all $\mathbf{x} \in \partial\mathcal{C}$ it holds that $\partial h(\mathbf{x})/\partial \mathbf{x} \neq 0$, then any Lipschitz continuous controller $\mathbf{k}(\mathbf{x}) \in \mathcal{U}_{\text{CBF}}(\mathbf{x})$ for the system in Equation 1 renders the set \mathcal{C} positively invariant.*

Intuitively, any controller $\mathbf{k}(\mathbf{x}) \in \mathcal{U}_{\text{CBF}}$ guarantees that the Lie derivative of h is nonnegative on the boundary of \mathcal{C} . Since h is only nonnegative on \mathcal{C} , the set is positively invariant under the assumption of locally Lipschitz closed-loop dynamics.

Any arbitrary given feedback controller might not satisfy the condition on the Lie derivative from Equation 5 of the control barrier function h . Freeman and Kototovic (1996); Ames et al.

(2014) introduce a point-wise min-norm controller that minimally modifies the given controller to guarantee positive control invariance. As the CBF condition on positive invariance is affine in the control input \mathbf{u} , we can define a quadratic program (QP) to find the minimum adjustment to a given controller $\mathbf{k}(\mathbf{x})$ that renders the controlled system positively invariant:

$$\mathbf{u}^*(\mathbf{x}) = \operatorname{argmin}_{\mathbf{u} \in \mathcal{U}} \frac{1}{2} \|\mathbf{u} - \mathbf{k}(\mathbf{x})\|_2^2 \quad (6a)$$

$$\text{s.t. } L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} \geq -\gamma(h(\mathbf{x})). \quad (6b)$$

Intuitively, the optimized $\mathbf{u}^*(\mathbf{x})$ guarantees that the set \mathcal{C} is positively invariant and stays as close as possible to the original controller $\mathbf{k}(\mathbf{x})$, where closeness is specified with respect to a chosen distance measure (e.g., the Euclidean norm in Equation 6a). Using this approach, we can apply any controller to the system, and the safety filter in Equation 6 guarantees positive invariance of the system on the safe set \mathcal{C} .

Consider the true system as a disturbed nominal system as in Equation 3, where $\mathbf{d}(\mathbf{x}, \mathbf{u}) \in \mathbb{D} \subset \mathbb{R}^n$ is the model mismatch between the nominal and the true system. The Lie derivative of the disturbed dynamics is given by

$$\frac{\partial h(\mathbf{x})}{\partial \mathbf{x}} \dot{\mathbf{x}} = L_{\hat{\mathbf{f}}}h(\mathbf{x}) + L_{\hat{\mathbf{g}}}h(\mathbf{x})\mathbf{u} + \delta(\mathbf{x}, \mathbf{u}) = L_{\hat{\mathbf{f}}_{\text{cl}}}h(\mathbf{x}) + \delta(\mathbf{x}, \mathbf{u}), \quad (7)$$

where $\delta(\mathbf{x}, \mathbf{u}) = \partial h(\mathbf{x})/\partial \mathbf{x} \mathbf{d}(\mathbf{x}, \mathbf{u}) \in \mathbb{R}$ is the error in the Lie derivative. We assume that $|\delta(\mathbf{x}, \mathbf{u})| \leq \rho$ for all $(\mathbf{x}, \mathbf{u}) \in \mathcal{X} \times \mathcal{U}$. For a state-feedback controller $\mathbf{u} = \mathbf{k}(\mathbf{x})$, the Lie derivative of the disturbed closed-loop system is given by $L_{\mathbf{f}_{\text{cl}}}h(\mathbf{x}) := L_{\hat{\mathbf{f}}_{\text{cl}}}h(\mathbf{x}) + \delta(\mathbf{x}, \mathbf{u})$. The notion of projection-to-state safety (PSSf) determines safety for uncertain systems. We use a similar definition compared to Taylor et al. (2021); however, we directly use the control barrier function candidate $h(\mathbf{x})$ as the projection:

Definition 5 (Projection-to-state safety (Taylor et al., 2021)) *The disturbed control system in Equation 3 is projection-to-state safe (PSSf) on $\mathcal{C}_{\bar{h}} = \{\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^n \mid \bar{h}(\mathbf{x}) \geq 0\}$ if there exists $\rho > 0$ such that the safe set $\mathcal{C} \supset \mathcal{C}_{\bar{h}}$ is positively control invariant for all $\delta(\mathbf{x}, \mathbf{u})$ satisfying $|\delta(\mathbf{x}, \mathbf{u})| \leq \rho$ for all $(\mathbf{x}, \mathbf{u}) \in \mathcal{X} \times \mathcal{U}$, where $\bar{h}(\mathbf{x}) = h(\mathbf{x}) - \rho$ with the control barrier function h as the projection and ρ being the bound of the projected disturbances δ in Equation 7.*

Under the assumption that the set $\mathcal{C}_{\bar{h}}$ is non-empty, we have the following lemma:

Lemma 6 *Let $\mathcal{C}_{\bar{h}} \subset \mathbb{R}^n$ be the 0-superlevel set of a C^1 function $\bar{h} : \mathbb{R}^n \rightarrow \mathbb{R}$ and for all $\mathbf{x} \in \partial \mathcal{C}_{\bar{h}}$ it holds that $\partial \bar{h}(\mathbf{x})/\partial \mathbf{x} \neq 0$. If there exists a Lipschitz continuous state-feedback controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that $L_{\mathbf{f}}\bar{h}(\mathbf{x}) + L_{\mathbf{g}}\bar{h}(\mathbf{x})\mathbf{k}(\mathbf{x}) \geq -\gamma(\bar{h}(\mathbf{x}))$, and there exists $\rho > 0$ such that the projected disturbances $\delta(\mathbf{x}, \mathbf{k}(\mathbf{x}))$ in Equation 7 satisfy $|\delta(\mathbf{x}, \mathbf{k}(\mathbf{x}))| \leq \rho$ for all $(\mathbf{x}, \mathbf{k}(\mathbf{x})) \in \mathcal{X} \times \mathcal{U}$ and $\rho \leq -\gamma(-\rho)$, then there exists $\gamma_{\text{PSSf}} \in \mathcal{K}_{\infty, e}$ that renders the disturbed system in Equation 3 PSSf on $\mathcal{C}_{\bar{h}}$.*

Proof Let $\gamma_{\text{PSSf}}(h(\mathbf{x})) = \gamma(h(\mathbf{x}) - \rho) - \gamma(-\rho)$, then $\gamma_{\text{PSSf}} \in \mathcal{K}_{\infty, e}$ since $\gamma_{\text{PSSf}}(0) = 0$ and

$$L_{\mathbf{f}_{\text{cl}}}(\bar{h}(\mathbf{x}) + \rho) = L_{\hat{\mathbf{f}}_{\text{cl}}}\bar{h}(\mathbf{x}) + \delta(\mathbf{x}, \mathbf{k}(\mathbf{x})) \quad (8a)$$

$$\geq L_{\hat{\mathbf{f}}_{\text{cl}}}\bar{h}(\mathbf{x}) - \rho \quad (8b)$$

$$\geq -\gamma(\bar{h}(\mathbf{x})) - \rho \quad (8c)$$

$$\begin{aligned} & \text{using } \rho \leq -\gamma(-\rho) \\ & \geq -\gamma(\bar{h}(\mathbf{x})) + \gamma(-\rho) \end{aligned} \quad (8d)$$

$$= -\gamma_{PSSf}(\bar{h}(\mathbf{x}) + \rho). \quad (8e)$$

Then, the disturbed system in Equation 3 is positively control invariant on \mathcal{C} and PSSf on $\mathcal{C}_{\bar{h}}$. \blacksquare

In the case that $\rho > -\gamma(-\rho)$, only the greater 0-superlevel set of $h(\mathbf{x}) + h_0$ can be rendered positively invariant instead of \mathcal{C} , where $h_0 = \gamma^{-1}(-\rho) + \rho$.

3.3. Bayesian Linear Regression

In this section, we briefly describe model learning using Bayesian linear regression (BLR), see [Bishop and Nasrabadi \(2006\)](#) for details. Our goal is to learn the Lie derivative residual as a linear combination of basis functions with additive Gaussian noise using a set of data points. Given a dataset $\mathcal{D} = \{\mathbf{X}, \mathbf{y}\}$ of independent identically distributed data points, where the input set is $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ with $\mathbf{x}_i \in \mathbb{R}^n$ and $N > 0$ and the target set is $\mathbf{y} = (y_1, \dots, y_N)$ with $y_i \in \mathbb{R}$. We assume that the Lie derivative residual can be modeled as $\mathbf{w}^\top \phi(\mathbf{x})$, where $\mathbf{w} \in \mathbb{R}^M$ with $M > 0$ are the weights and $\phi(\mathbf{x}) = (\phi_1(\mathbf{x}), \dots, \phi_M(\mathbf{x}))^\top$ defines the set of basis functions. We want to determine the weights such that the BLR model closely represents the dataset \mathcal{D} . We assume a Gaussian prior distribution over the weights \mathbf{w} as $p(\mathbf{w}) = \mathcal{N}(\mu_0, \Sigma_0)$, with mean μ_0 and covariance Σ_0 . Let $\Phi \in \mathbb{R}^{N \times M}$ be the design matrix with entries $\phi_{i,j} = \phi_j(\mathbf{x}_i)$ and σ^2 is the noise variance. Then, the posterior distribution is given by $p(\mathbf{w}|\mathcal{D}) = \mathcal{N}(\mu_N, \Sigma_N)$, where $\mu_N = \Sigma_N (\Sigma_0^{-1} \mu_0 + 1/\sigma^2 \Phi^\top \mathbf{y})$ and $\Sigma_N^{-1} = \Sigma_0^{-1} + 1/\sigma^2 \Phi^\top \Phi$. For sequentially arriving data, the previous posterior can be used as the prior for the following data point. The predictive distribution is given by $p(y|\mathbf{x}, \mathcal{D}) = \mathcal{N}(\mu_{\text{pred}}, \sigma_{\text{pred}}^2)$, where $\mu_{\text{pred}} = \mu_N^\top \phi(\mathbf{x})$ and $\sigma_{\text{pred}}^2 = \phi(\mathbf{x})^\top \Sigma_N \phi(\mathbf{x}) + 1/\sigma^2$.

4. Safely Learning Control Barrier Conditions

In this section, we present our barrier Bayesian linear regression (BBLR) approach that safely learns uncertain control barrier conditions online. By construction, we guarantee that (i) the Lie derivative residual learned by the BBLR is 0 on the boundary of the safe set $\partial\mathcal{C}$ and (ii) there exists an appropriate extended class \mathcal{K}_∞ function to lower bound the Lie derivative in the CBF condition.

4.1. Exploiting Control Barrier Condition Structure for Safe Online Learning

The model mismatch between the true and nominal systems directly impacts the Lie derivative in the CBF condition. The Lie derivative residual can be written as $\delta(\mathbf{x}, \mathbf{u}) = \alpha(\mathbf{x}) + \beta(\mathbf{x})^\top \mathbf{u}$, where $\alpha(\mathbf{x}) = \partial h(\mathbf{x})/\partial \mathbf{x} \mathbf{a}(\mathbf{x})$ and $\beta(\mathbf{x}) = \partial h(\mathbf{x})/\partial \mathbf{x} \mathbf{b}(\mathbf{x})$. We aim to safely learn an estimate $\hat{\delta}(\mathbf{x}, \mathbf{u}) = \hat{\alpha}(\mathbf{x}) + \hat{\beta}(\mathbf{x})^\top \mathbf{u}$ of the Lie derivative residual $\delta(\mathbf{x}, \mathbf{u})$ from online collected data. Our proposed online learned safety filter augments an arbitrary Lipschitz continuous controller $\mathbf{k}(\mathbf{x})$ using a QP (similar to Equation 6) to guarantee that the true system stays inside the safe set \mathcal{C} , while also expanding the set of safe control inputs $\mathcal{U}_{CBF}(\mathbf{x})$ for the true system. We exploit the structure of the CBF condition by choosing the set of basis functions as compositions of design functions ϕ_j and the control barrier function h . By construction, this will guarantee that the learned Lie derivative residual will have a value greater than or equal to zero on the boundary $\partial\mathcal{C}$. In contrast, arbitrary function approximators typically do not satisfy this condition.

Since $h(\mathbf{x})$ is C^1 and \mathcal{C} is compact, h attains a minimum $h_{\min} = 0$ and maximum $h_{\max} = \max_{\mathbf{x} \in \mathcal{C}} h(\mathbf{x})$ over \mathcal{C} , and the codomain of $h : \mathcal{C} \rightarrow \mathcal{H}$ with $\mathcal{H} = [0, h_{\max}]$ is compact. Let \mathcal{F} be the set of functions that are lower bounded by a negative class $\mathcal{K}_{\infty, e}$ function on \mathcal{H} with $\mathcal{F} = \{\phi : \mathcal{H} \rightarrow \mathbb{R} \mid \exists \gamma \in \mathcal{K}_{\infty, e}, \forall r \in \mathcal{H}, \phi(r) \geq -\gamma(r)\}$. Then, we have the following results:

Lemma 7 *The set \mathcal{F} is closed under addition.*

Proof Let $\phi_j \in \mathcal{F}$ for $j \in \{1, \dots, N\}$ be a finite collection of functions. By definition there exist γ_j such that for all $j \in \{1, \dots, N\}$, $\phi_j(r) \geq -\gamma_j(r)$. This implies that $\sum_{j=1}^N \phi_j(r) \geq -\sum_{j=1}^N \gamma_j(r) =: -\gamma(r)$, where we used the result from Kellett (2014), that a finite sum of class $\mathcal{K}_{\infty, e}$ functions is also of class $\mathcal{K}_{\infty, e}$. ■

Lemma 8 *The CBF condition for the estimated Lie derivative $L_{\hat{\mathbf{f}}_{\text{cl}}} \bar{h}(\mathbf{x}) + \hat{\alpha}(\mathbf{x}) + \hat{\beta}(\mathbf{x})^\top \mathbf{u} \geq -\gamma(h(\mathbf{x}))$ is satisfied, if there exist $\gamma_0, \gamma_\delta \in \mathcal{K}_{\infty, e}$ such that $L_{\hat{\mathbf{f}}_{\text{cl}}} \bar{h}(\mathbf{x}) \geq -\gamma_0(h(\mathbf{x}))$ and $\hat{\alpha}(\mathbf{x}) + \hat{\beta}(\mathbf{x})^\top \mathbf{u} \geq -\gamma_\delta(h(\mathbf{x}))$.*

Proof Using the result from (Kellett, 2014) on the summation of a finite number of class $\mathcal{K}_{\infty, e}$ functions, we have $L_{\hat{\mathbf{f}}_{\text{cl}}} \bar{h}(\mathbf{x}) + \hat{\alpha}(\mathbf{x}) + \hat{\beta}(\mathbf{x})^\top \mathbf{u} \geq -\gamma_0(h(\mathbf{x})) - \gamma_\delta(h(\mathbf{x})) =: -\gamma(h(\mathbf{x}))$. Therefore, the CBF condition in Equation 5 is satisfied. ■

The main result of our paper combines the previous Lemmas to guarantee safe control barrier conditions using linear combinations of basis functions $\phi_j \circ h$. The idea is to guarantee that any learned Lie derivative residual $\hat{\delta}$ is lower bounded by an extended class \mathcal{K}_{∞} function.

Theorem 9 *Let $L_{\hat{\mathbf{f}}_{\text{cl}}} \bar{h}(\mathbf{x}) + \hat{\alpha}(\mathbf{x}) + \hat{\beta}(\mathbf{x})^\top \mathbf{u}$ be the estimate of the true Lie derivative $L_{\mathbf{f}_{\text{cl}}} \bar{h}(\mathbf{x})$ and $\mathbf{u} \in \mathcal{U}$, where \mathcal{U} is compact, and $\hat{\alpha}$ and $\hat{\beta}$ are defined as*

$$\hat{\alpha}(\mathbf{x}) = \sum_{j=1}^N \lambda_j \phi_j(h(\mathbf{x})), \quad \hat{\beta}(\mathbf{x}) = \left(\sum_{j=1}^N \mu_{1,j} \phi_j(h(\mathbf{x})) \quad \dots \quad \sum_{j=1}^N \mu_{m,j} \phi_j(h(\mathbf{x})) \right)^\top \quad (9)$$

with nonnegative weights $\lambda_j, \mu_{i,j} \in \mathbb{R}_{\geq 0}$, $\phi_j \in \mathcal{F}$. If there exists a $\gamma_0 \in \mathcal{K}_{\infty, e}$ such that $L_{\hat{\mathbf{f}}_{\text{cl}}} \bar{h}(\mathbf{x}) \geq -\gamma_0(h(\mathbf{x}))$, then $L_{\hat{\mathbf{f}}_{\text{cl}}} \bar{h}(\mathbf{x}) + \hat{\alpha}(\mathbf{x}) + \hat{\beta}(\mathbf{x})^\top \mathbf{u} \geq -\gamma(h(\mathbf{x}))$ with $\gamma \in \mathcal{K}_{\infty, e}$.

Proof From Lemma 8, since $L_{\hat{\mathbf{f}}_{\text{cl}}} \bar{h}(\mathbf{x}) \geq -\gamma_0(h(\mathbf{x}))$ is satisfied by assumption, we need to show that $\hat{\alpha}(\mathbf{x}) + \hat{\beta}(\mathbf{x})^\top \mathbf{u} \geq -\gamma_\delta(h(\mathbf{x}))$ to satisfy the CBF condition. Equation 9 for all $\mathbf{x} \in \mathcal{C}_{\bar{h}}$ yields:

$$\hat{\alpha}(\mathbf{x}) + \hat{\beta}(\mathbf{x})^\top \mathbf{u} = \sum_{j=1}^N \lambda_j \phi_j(h(\mathbf{x})) + \sum_{i=1}^m \sum_{j=1}^N \mu_{i,j} u_i \phi_j(h(\mathbf{x})) \quad (10a)$$

$$\stackrel{\text{using } \phi_j \in \mathcal{F}}{\geq} - \sum_{j=1}^N \lambda_j \gamma_j(h(\mathbf{x})) - \sum_{i=1}^m \sum_{j=1}^N \mu_{i,j} |u_i| \gamma_j(h(\mathbf{x})) \quad (10b)$$

$$\stackrel{\text{using Lemma 7}}{=} -\gamma_\delta(h(\mathbf{x})), \quad (10c)$$

where $\gamma_\delta \in \mathcal{K}_{\infty, e}$. ■

Theorem 9 shows that any learned Lie derivative residual estimate that is a sum of nonnegatively weighted function compositions $\phi_j \circ h$ with $\phi_j \in \mathcal{F}$ is lower bounded by a negative class $\mathcal{K}_{\infty,e}$ function. This guarantees that the sum of the nominal Lie derivative $L_{f_{cl}} h(\mathbf{x})$ and the learned Lie derivative residual $(\hat{\alpha}(\mathbf{x}) + \hat{\beta}(\mathbf{x})^\top \mathbf{u})$ satisfies the CBF condition as well.

4.2. Locally Learning Control Barrier Conditions with Online Bayesian Linear Regression

A common assumption in previous work is that the learned Lie derivative satisfies the CBF condition. Additionally, the learned Lie derivative did not account for the uncertainty in the numerical Lie derivative estimation. We explicitly guarantee the satisfaction of the CBF condition, by leveraging the structure of our presented Lie derivative residual from the previous section in an online framework. We consider the uncertainty in the numerical Lie derivative estimates through Bayesian linear regression, which yields the uncertainty in our predictions.

Since the true system is unknown, we cannot determine the Lie derivative of the true system. Therefore, we estimate the Lie derivative residual locally from recent measurements summarized in a dataset $\mathcal{D}_{t_0,t_1} = \{\mathbf{X}_{t_0,t_1}, \mathbf{U}_{t_0,t_1}, \mathbf{H}_{t_0,t_1}\}$, where $\mathbf{X}_{t_0,t_1} = (\mathbf{x}_{t_0}, \dots, \mathbf{x}_t, \dots, \mathbf{x}_{t_1})$ with $\mathbf{x}_t = \mathbf{x}(t\Delta_t)$, $\mathbf{U}_{t_0,t_1} = (\mathbf{u}_{t_0}, \dots, \mathbf{u}_{t_1})$, and $\mathbf{H}_{t_0,t_1} = (h(\mathbf{x}_{t_0}), \dots, h(\mathbf{x}_{t_1}))$ where $t_1 \geq t \geq t_0 \geq 0$, $\Delta_t > 0$ is the sampling time, $t_1 = k - 1$, where k is the current discretized time step, $t_0 = t_1 - T$, and T is the number of past time steps we consider. Then, using the forward difference method we can determine an estimate of the true Lie derivative: $L_{f_{cl}} h(\mathbf{x}) \approx (h(\mathbf{x}_{t+1}) - h(\mathbf{x}_t))/\Delta_t$.

Remark 10 *The Lie derivative estimation error can be bounded following (Nocedal and Wright, 2017) with $c\Delta_t/2 := \sigma_{\text{diff}}$, where $c = \sup_{\mathbf{x} \in \mathcal{C}} d^2 h(\mathbf{x})/dt^2$. In practice, the second or any higher time derivative of h cannot be determined. Instead, Δ_t needs to be chosen sufficiently small and c needs to be chosen conservatively to account for the numerical differentiation errors.*

Our goal is to use BLR to determine the weights for the basis functions from these local datasets. We assume a Gaussian prior over the weights for BLR. This violates the condition of nonnegative weights from **Theorem 9**. To guarantee that every basis function still satisfies the conditions, we reduce the set of possible basis functions to functions for which $\phi_j \in \mathcal{F}$ and $-\phi_j \in \mathcal{F}$ on the non-negative domain \mathcal{H} , where \mathcal{F} is the class of functions that are lower bounded by negative class $\mathcal{K}_{\infty,e}$ functions. This restriction requires for all $\mathbf{x} \in \mathcal{C}$, $\phi_j(h(\mathbf{x})) = 0$. This yields safety, however, this excludes positive Lie derivative residuals on the boundary. The class of locally Lipschitz functions on the set \mathcal{H} for which $\phi_j(0) = 0$ satisfies this restriction. The following corollary guarantees safety on safe set \mathcal{C} with the more restricted set of basis functions:

Corollary 11 *Let $\hat{\delta}(\mathbf{x}, \mathbf{u}) = \mathbf{w}^\top \Phi(\mathbf{x}, \mathbf{u})$, with $\mathbf{w} \in \mathbb{R}^{N(m+1)}$ and $\Phi(\mathbf{x}, \mathbf{u}) = (\Phi_{\mathbf{x}}(\mathbf{x})^\top, \Phi_{\mathbf{u}}(\mathbf{x}, \mathbf{u})^\top)^\top$, where $\Phi_{\mathbf{x}}(\mathbf{x}) = (\phi_1(h(\mathbf{x})), \dots, \phi_N(h(\mathbf{x})))^\top$ and $\Phi_{\mathbf{u}}(\mathbf{x}, \mathbf{u}) = (u_1 \Phi_{\mathbf{x}}(\mathbf{x})^\top \dots u_m \Phi_{\mathbf{x}}(\mathbf{x})^\top)^\top$ and ϕ_j is such that $\phi_j \in \mathcal{F}$ and $-\phi_j \in \mathcal{F}$. Then, $L_{f_{cl}} \bar{h}(\mathbf{x}) + \hat{\alpha}(\mathbf{x}) + \hat{\beta}(\mathbf{x})^\top \mathbf{u} \geq -\gamma(h(\mathbf{x}))$ is satisfied, if there exists $\gamma_0 \in \mathcal{K}_{\infty,e}$ such that $L_{f_{cl}} \bar{h}(\mathbf{x}) \geq -\gamma_0(h(\mathbf{x}))$.*

Proof The estimated Lie derivative residual $\hat{\delta}(\mathbf{x}, \mathbf{u})$ satisfies

$$\hat{\delta}(\mathbf{x}, \mathbf{u}) = \sum_{j=1}^N w_j \phi_j(h(\mathbf{x})) + \sum_{i=1}^m \sum_{j=1}^N w_{iN+j} u_i \phi_j(h(\mathbf{x})), \quad (11a)$$

$$\geq - \sum_{j=1}^N |w_j| \gamma_j(h(\mathbf{x})) - \sum_{i=1}^m \sum_{j=1}^N |w_{iN+j}| \|\mathbf{u}\|_{\infty} \gamma_j(h(\mathbf{x})), \quad (11b)$$

$$=: -\gamma_{\hat{\delta}}(h(\mathbf{x})). \quad (11c)$$

The rest of the proof follows similar arguments as the proof for [Theorem 9](#). ■

The restricted set of basis functions satisfies the safe application of BLR, which yields the predicted mean μ_k and covariance Σ_k . The estimated Lie derivative residual is $\hat{\delta}_k(\mathbf{x}, \mathbf{u}) = \mu_k^{\top} \Phi(\mathbf{x}, \mathbf{u})$ and the covariance of the Lie derivative residual is given by $\sigma_k(\mathbf{x}, \mathbf{u}) = \Phi(\mathbf{x}, \mathbf{u})^{\top} \Sigma_k \Phi(\mathbf{x}, \mathbf{u}) + \frac{1}{\sigma_{\text{diff}}^2}$. We assume that the resulting probabilistic model is well calibrated, such that there exists $s \in \mathbb{N}$ so that $\delta(\mathbf{x}, \mathbf{u}) \in \hat{\delta}_k(\mathbf{x}, \mathbf{u}) \pm s\sigma_k(\mathbf{x}, \mathbf{u})$ with high probability. Then, we can use our Lie derivative estimate to render an arbitrary Lipschitz continuous controller $\mathbf{k}(\mathbf{x})$ safe using the following QP:

$$\begin{aligned} \mathbf{u}_{\text{bblr}}^*(\mathbf{x}) = \operatorname{argmin}_{\mathbf{u} \in \mathcal{U}} & \quad \frac{1}{2} \|\mathbf{u} - \mathbf{k}(\mathbf{x})\|_2^2 \\ \text{s.t.} & \quad L_{\hat{f}} \bar{h}(\mathbf{x}) + L_{\hat{g}} \bar{h}(\mathbf{x}) \mathbf{u} + \hat{\delta}(\mathbf{x}, \mathbf{u}) - \rho \geq -\gamma_{PSSf}(\bar{h}(\mathbf{x}) + \rho) - \gamma_{\hat{\delta}}(h(\mathbf{x})), \end{aligned} \quad (12)$$

where the CBF condition in [Equation 12](#) follows by combining the results from [Corollary 11](#), and [Lemma 6](#) with [Equation 8b](#) for the left hand side, and [Equation 8e](#) for the right hand side of the constraint. Finally, we account for the uncertainty between the nominal Lie derivative and the true Lie derivative, which is captured by the constant ρ , and the uncertainty between the learned Lie derivative and true Lie derivative, which is captured by $\sigma_k(\mathbf{x}, \mathbf{u})$. Depending on these uncertainties, we blend the safe control inputs from a nominal filter $\mathbf{u}_{\text{nominal}}^*(\mathbf{x})$, which is a QP filter subject to the CBF condition as in [Lemma 6](#), and our proposed filter, which uses $\mathbf{u}_{\text{bblr}}^*(\mathbf{x})$ from [Equation 12](#). This yields $\mathbf{u}^* = (1-r)\mathbf{u}_{\text{nominal}}^*(\mathbf{x}) + r\mathbf{u}_{\text{bblr}}^*(\mathbf{x})$, where $r = q\left(\frac{\rho}{\rho + s\sigma_k(\mathbf{x}, \mathbf{u})}\right)$ with a monotonically increasing function q bounded by 1, and $s \in \mathbb{N}$, which determines the number of considered standard deviations. In contrast to previous work ([Taylor et al., 2021](#)), which blended between unsafe and safe control inputs during learning, our approach blends two safe control inputs and guarantees that the system will not leave \mathcal{C} during learning. When the uncertainty of the learned Lie derivative residual is larger than ρ , the system returns to the smaller safe set $\mathcal{C}_{\bar{h}}$ instead, due to the set's attractiveness ([Ames et al., 2019](#)), until the uncertainty is reduced again.

5. Simulation Results

We show the performance of our proposed BBLR on the example of an inverted pendulum:

$$\dot{x}_1 = x_2, \quad \dot{x}_2 = -g/l \sin(x_1) + 1/(ml^2) u, \quad (13)$$

where $(x_1, x_2)^{\top} = (\theta \quad \dot{\theta})^{\top}$ with θ being the angle of the pendulum from the downward position and $\dot{\theta}$ being the angular velocity, g is the gravitational constant, l is the length and m is the mass of the pendulum, and u is the input torque. The parameters of the true and nominal systems are $(l = 1.0 \text{ m}, m = 1.0 \text{ kg})$ and $(l = 1.0 \text{ m}, m = 0.96 \text{ kg})$, respectively. We conservatively estimate $\rho = 0.5$. We are given the CBF $h(\theta, \dot{\theta}) = 1 - (\theta/\theta_{\max})^2 - (\dot{\theta}/\dot{\theta}_{\max})^2$, where $\theta_{\max} = \pi \text{ rad}$ and $\dot{\theta}_{\max} = 5.0 \text{ rad/sec}$, which yields an ellipsoidal safe set that prevents the pendulum from swinging beyond the upright position. We design a controller based on the nominal system with the goal of

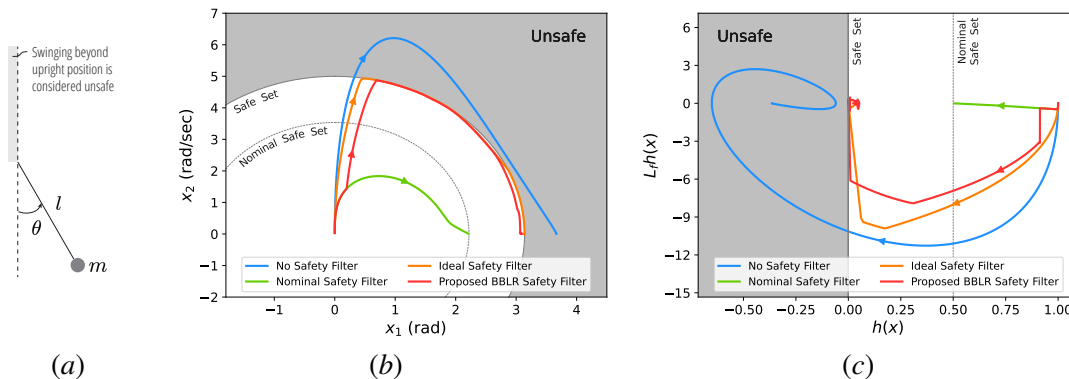


Figure 2: (a) The schematic of the pendulum; (b) the state trajectories demonstrate that our proposed BBLR safe filter achieves similar performance as compared to the filter using the true CBF condition, while also guaranteeing safety of the safe set; (c) the Lie derivative over the CBF evaluated for the true system confirms that our learned Lie derivative is guaranteed to be 0 on the boundary of the safe set.

swinging up the pendulum (see Figure 2 (a)): $u = (-18.8 (\theta - \pi) - 6.1 \dot{\theta} + 5)$ Nm. For BBLR we consider the set of monomials from degree 1 to 5 composited with $h(\theta, \dot{\theta})$ as the basis and learn from the last 0.2 s of data. For the first 0.2 s, we use the nominal filter and afterwards we apply the blended safe inputs as described in the previous section.

Figure 2 (b) shows the successful application of the learned CBF filter using BBLR. The unfiltered controller, leaves the safe set, while the filtered controllers keep the system inside the safe set. Our learned safety filter is able to keep the system at the boundary of the safe set, whereas the nominal filter conservatively keeps the system inside a conservative nominal safe set. Finally, Figure 2 (c) demonstrates the satisfaction of the CBF condition through BBLR as the Lie derivative evaluated for the true system reduces to 0 at the boundary. In the case of unfiltered control inputs, the system achieves a negative Lie derivative on the boundary resulting in the violation of safety. The nominal filter restricts the Lie derivative to 0 at the boundary of the more conservative nominal safe set.

6. Conclusions

We proposed the BBLR approach for safely learning control barrier conditions online. Our approach can be applied to partially unknown nonlinear control-affine systems with bounded uncertainty. By leveraging a special set of basis functions that are lower-bounded by negative class $\mathcal{K}_{\infty, e}$ functions, we were able to guarantee the satisfaction of the CBF condition while learning the Lie derivative residual using BLR. We exploit the predictive uncertainty provided by the BLR to systematically account for the errors in the online learning scheme. We demonstrated the improved performance and safety of our approach on a simulated two-dimensional pendulum.

References

- Aaron D. Ames, Jessy W. Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs with application to adaptive cruise control. *Proceedings of the IEEE Conference on Decision and Control (CDC)*, pages 6271–6278, 2014.
- Aaron D. Ames, Xiangru Xu, Jessy W. Grizzle, and Paulo Tabuada. Control Barrier Function Based Quadratic Programs for Safety Critical Systems. *IEEE Transactions on Automatic Control*, 62(8): 3861–3876, 2017.
- Aaron D. Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *Proceedings of the European Control Conference (ECC)*, pages 3420–3431, 2019.
- Christopher M. Bishop and Nasser M. Nasrabadi. *Pattern recognition and machine learning*. Springer, 2006.
- Lukas Brunke, Melissa Greeff, Adam W. Hall, Zhaocong Yuan, Siqi Zhou, Jacopo Panerati, and Angela P. Schoellig. Safe learning in robotics: From learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems*, 2022. In press, available at <https://arxiv.org/abs/2108.06266>.
- Richard Cheng, Gábor Orosz, Richard M. Murray, and Joel W. Burdick. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33(1), pages 3387–3395, 2019.
- Jason Choi, Fernando Castañeda, Claire Tomlin, and Koushil Sreenath. Reinforcement learning for safety-critical control under model uncertainty, using control lyapunov functions and control barrier functions. In *Proceedings of the Robotics: Science and Systems (RSS) Conference*, 2020.
- Jason Choi, Donggun Lee, Koushil Sreenath, Claire J. Tomlin, and Sylvia L. Herbert. Robust control barrier–value functions for safety-critical control. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, pages 6814–6821, 2021.
- David D. Fan, Jennifer Nguyen, Rohan Thakker, Nikhilesh Alatur, Ali-akbar Agha-mohammadi, and Evangelos A. Theodorou. Bayesian learning-based adaptive control for safety critical systems. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 4093–4099, 2020.
- Randy A. Freeman and Petar V. Kototovic. *Robust Nonlinear Control Design: State-Space and Lyapunov Techniques*. Birkhauser Boston Inc., 1996.
- Christopher M. Kellett. A compendium of comparison function results. *Mathematics of Control, Signals, and Systems*, 26(3):339–374, 2014.
- Hassan K. Khalil. *Nonlinear Systems*. Pearson Education. Prentice Hall, 2002.
- Mohammad J. Khojasteh, Vikas Dhiman, Massimo Franceschetti, and Nikolay Atanasov. Probabilistic safety constraints for learned high relative degree system dynamics. In *Proceedings of the Conference on Learning for Dynamics and Control (LADC)*, volume 120 of *Proceedings of Machine Learning Research (PMLR)*, pages 781–792, 2020.

- Brett T. Lopez, Jean-Jacques E. Slotine, and Jonathan P. How. Robust adaptive control barrier functions: An adaptive and data-driven approach to safety. *IEEE Control Systems Letters (L-CSS)*, 5(3):1031–1036, 2021.
- Ian M. Mitchell, Alexandre M. Bayen, and Claire J. Tomlin. A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on Automatic Control*, 50(7):947–957, 2005.
- Mitio Nagumo. Über die Lage der Integralkurven gewöhnlicher Differentialgleichungen. *Proceedings of the Physico-Mathematical Society of Japan, 3rd Series*, 24:551–559, 1942.
- Jorge Nocedal and Stephen Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, 2017.
- Motoya Ohnishi, Li Wang, Gennaro Notomista, and Magnus Egerstedt. Barrier-certified adaptive reinforcement learning with applications to brushbot navigation. *IEEE Transactions on Robotics*, 35(5):1186–1205, 2019.
- Andrew Taylor, Andrew Singletary, Yisong Yue, and Aaron Ames. Learning for safety-critical control with control barrier functions. In *Proceedings of the Conference on Learning for Dynamics and Control (LADC)*, volume 120 of *Proceedings of Machine Learning Research (PMLR)*, pages 708–717, 2020.
- Andrew J. Taylor and Aaron D. Ames. Adaptive safety with control barrier functions. In *Proceedings of the IEEE American Control Conference (ACC)*, pages 1399–1405, 2020.
- Andrew J. Taylor, Victor D. D. Dorobantu, Hoang M. Le, Yisong Yue, and Aaron D. Ames. Episodic learning with control Lyapunov functions for uncertain robotic systems. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6878–6884, 2019.
- Andrew J. Taylor, Andrew Singletary, Yisong Yue, and Aaron D. Ames. A control barrier perspective on episodic learning via projection-to-state safety. *IEEE Control Systems Letters*, 5(3):1019–1024, 2021.
- Li Wang, Evangelos A. Theodorou, and Magnus Egerstedt. Safe learning of quadrotor dynamics using barrier certificates. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2460–2465, 2018.